

COMMENTARY

Disparate treatment and adverse impact in applied attrition modeling

Christopher M. Castille* and Ann-Marie R. Castille

Nicholls State University

*Corresponding author. Email: christopher.castille@nicholls.edu

Speer, Dutta, Chen, and Trussell (2019) provide an excellent overview of key practices on applied attrition modeling. With our commentary, we wish to elaborate on a decision point Speer and colleagues left open in the development of attrition models, namely the decision to examine protected classification information. Our contribution seems particularly relevant given popular press discussions regarding discriminatory employment activities enacted via artificial intelligence. For instance, Amazon developed an algorithm with the purpose of recruiting job candidates with the highest potential. Unfortunately, the process resulted in a bias against female candidates that could not be remedied in a timely fashion, resulting in leaders terminating the project (Dastin, 2018). Our concern is that a similar outcome might result here if biases against protected classes are not examined thoroughly.

Unfortunately, while there is value in studying protected classes and their relation to turnover, we have observed that legal teams might resist people analytics teams' efforts to examine protected classes in projects such as the development of attrition models, and so we hope to speak to those practitioners who are facing such an obstacle. Therefore, with our commentary, we call attention to what has in our observation been a problem in practice: gaining permission to analyze protected class information on employees in building attrition models. Drawing upon the adverse impact and disparate treatment literature, we highlight how both including and failing to acknowledge the role of protected class information can introduce legal exposure to the organization in question. Our key contribution involves clarifying how *analytical* ignorance of protected class information might increase an organization's legal exposure. We emphasize analytical ignorance here as meaning that the attrition modelers remain agnostic to protected class information and enact discriminatory policy in an illegal fashion. We hope to augment guidance provided by Speer et al. (2019) by equipping industrial and organizational (I-O) psychologists who are engaged in modeling attrition with steps to take to ensure that their actions are both in compliance with employment law and also create business value.

The legal risks of examining (or failing to examine) protected class information in applied attrition modeling

As Speer et al. (2019) noted, attrition modeling involves using available organizational data to estimate the probability of employee turnover. Such estimates in turn feed organizational decision making and workforce planning (e.g., hiring, retention initiatives, changes in compensation, promotion, etc.). For the sake of discussion, we will assume that attrition modelers hope to build a model that would trigger an employment decision (e.g., "high risk" individuals would be targeted for a discussion regarding a change in compensation, benefits, or some aspect of the employment arrangement). In other words, attrition research informs employment decision making as a

manner of policy. One can conceive of a manager who will be held responsible for optimizing turnover for the benefit of the organization. Armed with models that estimate the probability of exit for an individual employee or group of employees, this manager would then take steps (e.g., improve compensation or benefits, etc.) that are aimed at increasing voluntary forms of functional turnover and or decreasing avoidable forms of dysfunctional turnover.

To bring clarity to the issue of whether or not protected classification information should be examined, we draw upon the existing disparate treatment and adverse impact literature and provide corresponding illustrations that might help attrition modelers explain why protected class factors should be examined when modeling attrition. We would suggest that attrition modelers may in certain circumstances have a legal obligation to examine the role of protected classification information in the analysis of attrition so to avoid engaging in employment-related decisions that would constitute illegal discrimination. Such inclusions could be direct and include protected class factors (e.g., including gender and ethnicity into an attrition model) or indirect via including or creating predictors that are correlated with protected class factors (e.g., job class, postal codes). Although examining protected class information has the obvious business value of improving variance explained in turnover, thus improving our ability to optimize turnover, looking at protected class information may in practice be viewed in a more precarious manner, particularly by legal teams.

When building an attrition model, *disparate treatment* would be a concern if there was a clear intention to discriminate against protected classes. As this is the more obvious form of injustice, we believe that legal teams' sensitivity to this issue may motivate their reason for denying access to protected class factors such as employee sex and race. If the in-practice attrition model were to contain protected class information because of statistically significant associations with attrition outcomes, such inclusion would be viewed as circumstantial evidence of disparate treatment (see *Schwager v. Sun Oil Co. of Pa.*, 1979, p. 34). This could even be argued if the model included a factor that was contingent upon "race," such as in-group diversity. When put into practice, such an attrition model would treat affected groups differently based on protected classifications, constituting disparate treatment.

By contrast, *adverse impact* would be a concern if the same standards or procedures were applied to all individuals who interact with the organization but produce a substantial difference in employment-related outcomes. Adverse impact would occur if the developed attrition model, though not including protected class information, nevertheless produced outcomes that were associated with protected class factors (e.g., enhancements in compensation or benefits). A simple example involves an attrition model encouraging managers to allow African Americans to exit at a disproportionate rate relative to other classes for reasons that are correlated in one's population in a unique manner with race but also unrelated to performance in the job (e.g., commute length). We think the risk of adverse impact is particularly salient when attrition modeling takes advantage of big data sets and machine learning (i.e., data sets that contain large swaths of information, such as candidates' residential postal code; for related discussions, see Guenole, Ferrar, & Feinzig, 2017; O'Neill, 2016). For instance, including employees' postal codes in an analysis may tap into socioeconomic status, which can be a proxy for race or ethnicity (see Guenole et al., 2017). Commute length, which could be a predictor of attrition for the organization in question, may be computed using postal codes. Therefore, including commute length, a variable that at face value should not cause adverse impact, may produce a statistically biased outcome along protected class lines (Guenole et al., 2017). As should be evident, these effects are more subtle in their manifestation and, therefore, deserving of greater concern by attrition modelers using big data and the legal teams with which they work. Legal representatives may be less sensitive to adverse impact in this context, which is where I-O psychologists can play a valuable role.

This brings us to our key point: Attrition modelers may be legally obliged to examine whether employing attrition models would cause *adverse impact*. In our observation, legal teams may resist the examination of protected class characteristics because such examinations, if they are to

unearth discrimination in the organization, could place the organization at legal risk during a possible discoverability process. As we hope we have made abundantly clear, such an action may not be in accordance with employment law. Furthermore, we would venture to suggest that a legal team's refusal to have protected class information examined could itself be discoverable and raise questions regarding intent to discriminate (which an organization may not possess). I-O psychologists may, therefore, need to make their legal teams aware of these concerns during the attrition modeling process.

A simple three-step process for examining adverse impact in attrition models

Seeking to augment the guidance provided by Speer et al. (2019), as well as the guidance provided by other commentators, we have provided a process that attrition modelers can use to ensure their models comply with employment law. Additionally, HR professionals, particularly those with minimal analytics comprehension, might adapt this process (see the corresponding footnotes) to leverage the benefits of attrition modeling in creating business value while also ensuring compliance with employment law. Such individuals would not likely develop the models but need to have those models audited.

In the first step, modelers should include protected class factors in their attrition-model-building phase. In other words, all protected class factors should be present in the initial model development data set. Here, it is important to flag common correlates of protected class factors and attrition. Such correlations can be helpful in diagnosing the cause (or causes) of any protected class information–attrition associations that appear later in the analysis. If modelers do not include protected class information, an explanation should be furnished as to why this is the case and how adverse impact is avoided. Last, an attrition model that does not contain the protected class factors should be built, estimates of the probability of exit estimated, and decision rules crafted regarding whether a policy will be triggered (e.g., flagged individuals require job redesign). Ensuring that the model does not contain protected class factors essentially protects the organization in question from engaging in disparate treatment and is a necessary but not sufficient step.

In the second step, these decision rules are tested for possible adverse impact. Here, the decision rules (e.g., “1” means the individual is targeted for intervention) are examined via statistical test (e.g., chi-square tests of association, Fisher's exact test) with protected class factors. This test will be used to identify whether the associated employment decision is statistically associated with protected classifications. Here, a statistically significant association would suggest that adverse impact could be occurring.

In the third step, business leaders should act on the results of the adverse impact assessment. If there is no statistical evidence of adverse impact, then the attrition model would probably be fine for practical use. If there is statistical evidence of adverse impact, then the cause or causes of such impact should be further studied using accepted frameworks for examining bias to determine if unintended and illegal discrimination is occurring (see Guenole, 2018, for a discussion of this issue). Indeed, the appearance of adverse impact may not reflect bias against protected classes but may unearth true differences (even psychological ones) in the populations that are being studied. For instance, there is evidence of sex differences in vocational interests (e.g., men prefer more realistic occupations; see Su, Rounds, & Armstrong, 2009). There is also evidence that when individuals' interests align with their choice of occupation, they tend to hold lower turnover intentions and so are less likely to leave (Van Iddekinge, Roth, Putka, & Lanivich, 2011). Therefore, it is plausible that any correlation between sex and attrition within a given job or occupational context may be better explained by theories of person–job fit (e.g., attraction-selection-attrition; see Schneider, 1987) rather than unintentional or illicit discrimination. Such circumstantial covariation could give rise to the appearance of adverse impact where there is none. Unfortunately, attrition modelers would not realize this without conducting the appropriate tests and research.

Last, because the organization would like to capitalize on the benefits of attrition modeling while the cause(s) of adverse impact is being investigated further (or even if they are not being investigated further), they can simply remove all factors from the attrition model that give rise to a statistically significant assessment of adverse impact just to be safe. Although this will reduce the utility of the model that is used, this step should allow the organization to reap the benefits of attrition modeling without violating employment law.

Conclusion

We think it would be legally unwise for attrition modelers to ignore protected class information in their work. As attrition modeling is, indeed, a relatively quick win for an analytics team, there could be a pressure to play fast and loose with an organization's data. I-O psychologists, perhaps unlike the data scientists with whom they may work, should be proactive in ensuring that employment laws are not violated in delivering a quick win. If an organization has expressed a commitment to diversity and inclusion, then leveraging this prior commitment may help convince business leaders that protected class information should be considered at the outset of one's work. Indeed, three field studies and an experiment by Mayer, Ong, Sonenshein, and Ashford (2019) suggest that when organizations have expressed a commitment to diversity and inclusion, obliging moral action (i.e., asking what an organization with a commitment to diversity and inclusion should do) could help sell these issues to business leaders.

References

- Dastin, J.** (2018, October 9). Amazon scraps secret AI recruiting tool that showed bias against women. Retrieved from <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>
- Guenole, N.** (2018, July). Better the devil you know? Opportunities and risks with 21st century assessment. Presented at the CognitionX, Goldsmiths, University of London.
- Guenole, N., Ferrar, J., & Feinzig, S.** (2017). *The power of people: Learn how successful organizations use workforce analytics to improve business performance*. Indianapolis, IN: Cisco Press.
- Mayer, D. M., Ong, M., Sonenshein, S., & Ashford, S. J.** (2019). The money or the morals? When moral language is more effective for selling social issues. *Journal of Applied Psychology*, **104**(8), 1058–1076.
- O'Neil, C.** (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York, NY: Crown Publishing Group.
- Schneider, B.** (1987). The people make the place. *Personnel Psychology*, **40**(3), 437–453. doi: [10.1111/j.1744-6570.1987.tb00609.x](https://doi.org/10.1111/j.1744-6570.1987.tb00609.x).
- Schwager v. Sun Oil Co. of Pa.*, 591 F.2d 58 (10th Cir. 1979).
- Speer, A. B., Dutta S., Chen, M., & Trussell, G.** (2019). Here to stay or go? Connecting turnover research to applied attrition modeling. *Industrial and Organizational Psychology: Perspectives on Science and Practice*, **12**(3), 277–301.
- Su, R., Rounds, J., & Armstrong, P. I.** (2009). Men and things, women and people: A meta-analysis of sex differences in interests. *Psychological Bulletin*, **135**(6), 859–884. doi: [10.1037/a0017364](https://doi.org/10.1037/a0017364)
- Van Iddekinge, C. H., Roth, P. L., Putka, D. J., & Lanivich, S. E.** (2011). Are you interested? A meta-analysis of relations between vocational interests and employee performance and turnover. *Journal of Applied Psychology*, **96**(6), 1167–1194. doi: [10.1037/a0024343](https://doi.org/10.1037/a0024343)

Cite this article: Castille C.M. and Castille A.-M.R. (2019). Disparate treatment and adverse impact in applied attrition modeling. *Industrial and Organizational Psychology* **12**, 310–313. <https://doi.org/10.1017/iop.2019.53>